Agent-based Modelling of Command-and-Control Effectiveness

Michael F. Ling

Theatre Operations Branch

Aeronautical and Maritime Research Laboratory, DSTO

Abstract

Uncertainty is one of the most pervasive elements in war, and it manifests itself no more prominently than in the area of Command and Control (C2) where the human factor plays a crucial role. For example, the situation facing the commanders in the battlefield could be very different from that perceived at the HQ. In the dynamic environment of a military operation, uncertainty is highly difficult, if not impossible, to control because of strong human interactions. New behaviours could emerge out of these complex interactions. In this work, a high level agent-based model (ABM) is used to investigate how such uncertainties may hinder C2 operational effectiveness, even when high quality information is available. We also aim to understand some of the patterns of behaviour that may emerge from the complex interactions between the agents.

Introduction

Over the past couple of decades there has been an increasing realisation [1,2] that conventional operations research tools based on rigorous mathematical equations and detailed physical description of combat cannot provide a realistic description of the complex and dynamic situations in which military operations are conducted. In such operations, may they be warfighting or peacekeeping, the participants have to interact with hostile or potentially hostile forces, and to respond to the hostile actions. In the process, a new situation or environment is created, which in turn triggers off new responses from both sides. Furthermore, war arouses some of the strongest human emotions, which make it even more difficult to anticipate the behaviours of individuals in a command and control chain [3].

The main drawback of equation-based models (EBM) is that they are incapable of dealing with the dynamics of interactions between the combating sides and their reactions to each other's actions. Another serious challenge to EBM is that the world is fundamentally nonlinear, and consequently many problems defy the traditional scientific approach of analysis by decomposition. The nonlinearity of warfighting means that small changes in certain critical (initial) conditions can profoundly alter the outcomes, thus making reliable prediction extremely difficult, if not impossible. Finally, EBM cannot even begin to model the so-called intangibles such as human emotions, aggressiveness, fear, anger, team cohesion and trust [1].

With the advent of complexity theory and its application to warfare studies, a case has been made for viewing warfare as a complex adaptive system (CAS), which adapts, evolves and coevolves with its environment [2,4]. It follows that a C2 system should likewise be treated as a CAS, which involves interactions between different levels of commanders, and the commanders with the enemy. This dynamic and interactive process makes it quite impossible to achieve certainty at any time [5]. It also renders any serious attempt to study the C2 operation using EBM extremely difficult due to the dynamic and nonlinear nature of the problem.

Agent-based models (ABM) offer an opportunity to analyse the aforementioned complexity problems by concentrating on the behaviour of and interactions between the participating entities instead of the performance of specific weapons or sensors [2,4,6]. In other words, we shift our attention from analysing the performance of pieces of equipment to how different modes of operation may alter the outcome of a combat or peacekeeping or how the C2 system utilises information and acts upon it. In this approach, we concentrate on the emergent patterns of the *whole* (*e.g.* the fitness landscape) rather than the individual *parts* (*e.g.* the attrition rate).

Some of the notable examples of using intelligent agents to study emergent behaviour in warfare are the *Irreducible Semi-Autonomous Adaptive Combat* (ISAAC) and the *Enhanced ISAAC Neural Simulation Toolkit* (EINSTein), from the US Marine Corps Combat Development Command (MCCDC) as part of their *Project Albert* research [1,4,6]. Another useful model is the *Map Aware Non-uniform Automata* (MANA) [7,8] developed primarily with peacekeeping in mind by the Defence Operational Technology Support Establishment (DOTSE), New Zealand.

The Model

This work makes use of the *Map Aware Non-uniform Automata* (MANA) developed in DOTSE. MANA was based heavily on Ilachinski's cellular automaton land combat model ISAAC. The details of ISAAC have been discussed extensively by Ilachinski [4,6] and are available from the Center for Naval Analysis Website. Briefly, ISAAC models two competing teams of agents that each reacts to their surroundings guided by local information and an internal rule set. The model contains approximately thirty parameters governing the personality of each agent, weapon and sensor range and firepower, communication range between agents and squads. Each run of the model is stochastic and it produces several measures of fitness for each team.

MANA has fewer parameters than ISAAC [8] but also introduces two highly useful features. One is the trigger point, which enables the analyst to model the change of behaviour under a wider range of circumstances than allowed in ISAAC. The other is the situational awareness (SA) map generated from the head squad using its sensors. It is a representation of how good the information about the hostile forces is, and it is available to all agents on the same side though they may not necessarily use it to their full advantage. It is this capacity of the SA map that will be used in this work for studying the relationship between information superiority, its interpretation, uncertainty in C2, and mission effectiveness. It is worth bearing in mind that terms like combat and attrition in this work play merely the role of measure of fitness (MOF) for the performance of a team of agents.

The Baseline Scenario

The baseline scenario was set up with intercept-combat missions in mind. A squad of Blue agents is given two tasks: to capture the Red Post (Blues' waypoint) within a given time period and to intercept any Red agents on route in order to prevent them from capturing a Blue Post (Red's waypoint). These two tasks are given equal weight so that the Blues are not required to carry out an active search and destroy mission. They are only expected to intercept the Reds if they come within Blue's detection range (the Blues are supposed to have superior sensors and weapons). The MOF, measure of fitness, therefore takes into account the time taking to reach the Red Post, the number of Blue casualties and whether the Blue Post is taken by the Red agents.

The baseline scenario begins with a Red Squad of 8 with relatively good training, firepower, fire and sensor ranges, as well as good SA. On the Blue side, there are 18 elite agents, which are superior to the Red Squad in every respect. Furthermore, the Blue agents are given superior information by their UAV (unmanned aerial vehicle), who can also be considered as the Blues' HQ in this work, has the highest situational awareness as well as the greatest sensor range.

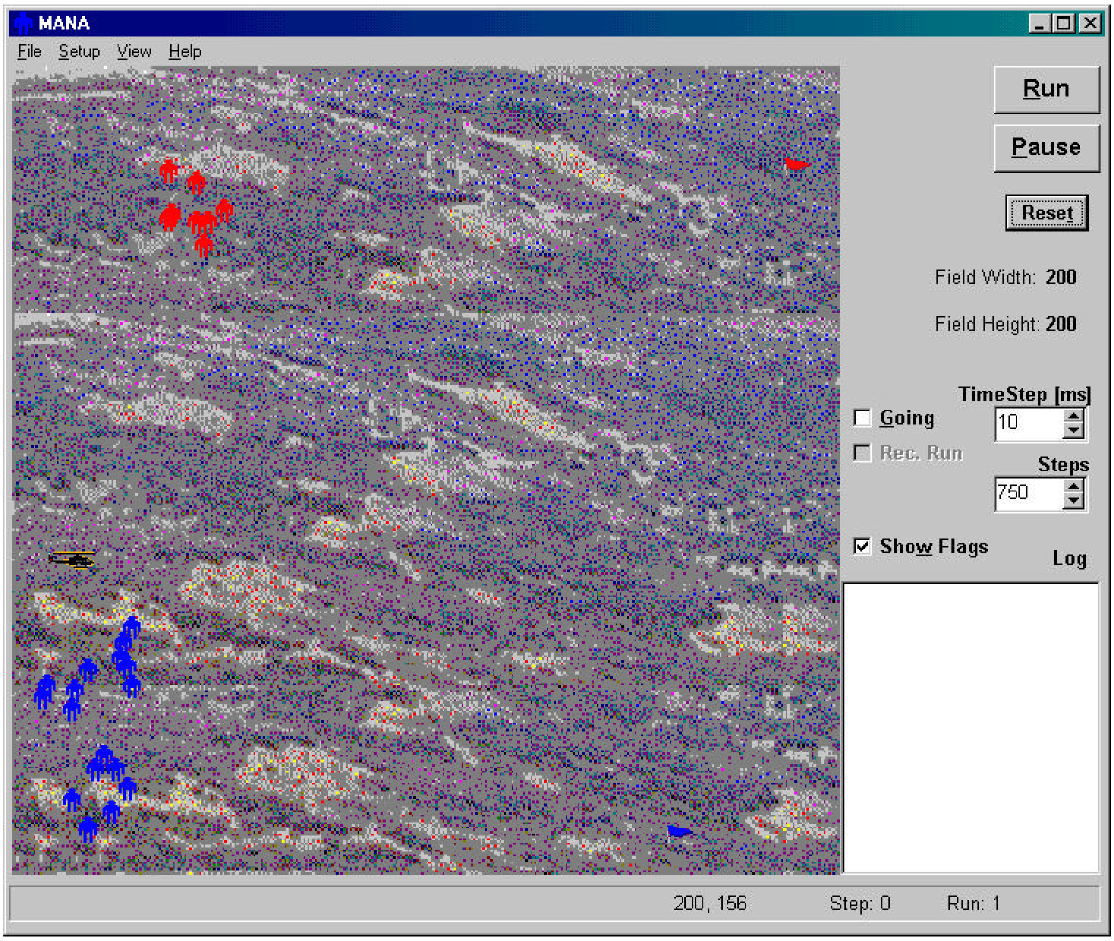

Figure 1 A typical MANA landscape with Blue and Red squads. The red and blue flags are the Red (Blue) and Blue (Red) Posts (waypoints), respectively. The helicopter symbol represents our UAV.

The Experiment

In this study we are concerned with two variables of the Blues: the value of situational awareness and the number of squads. The Blues' SA value can be considered as their ability to interpret the superior information passed on from the UAV, whereas the number of squads on the Blue side represents the number of sub-commands. The former gives a measure of how well each Blue squad makes use of the information for decision-making in its local environment. The latter introduces extra randomness (uncertainty) to the way the Blues handle their tasks. It should be noted that the actions of the Blue agents are meant to represent primarily the functioning of a C2 system, and not necessarily military actions as such.

In MANA, each side has a head squad, which, in accordance with its sensor range, set up a situational awareness map for its agents to use. In our study the head squad is the UAV (unmanned aerial vehicle) who has both the highest situational awareness and the greatest sensor range. The UAV is also considered the Blues' HQ purely for our C2 modelling purpose. Then the squads of Blue agents play the role of sub-commands of equal ranking, though they do not actually obey orders from the HQ (UAV). The idea is to observe the Blues' fitness for the tasks as different SA values are assigned to the squads. While the current version of MANA restricts us to effectively only two levels of command, it is possible to emulate a C2 system using the squad structure, and to investigate how the increase in squad (sub-command) numbers may complicate the interactions between the agents. It should be born in mind that every agent has a degree of randomness (entropy) in its movement and decision-making process, thus making it possible to simulate the (unpredictable) actions and responses of human individuals with different perceptions of the same situation.

Throughout this study, the Red Squad has 8 members with an SA value of 67%. All the relevant parameters governing their performance, such as firepower, weapon range, sensor range, and stealthiness are comparable to, but not as good as, the corresponding values of the elite Blue Squad. In the case of the Blues, only the number of squads and their SA values will be altered. The situational awareness of the UAV is chosen to 100% for all the scenarios in this work.

In the first experiment, an SA value of 86% is chosen for the Elite Blue Squad for it would be unrealistic to expect even elite agents to be able to make perfect use of available information. The Blue Squad is sent from the lower left-hand corner of the MANA map to capture the Red Post on the upper right-hand side. Almost diagonally cutting across their path, a Red Squad moves from the upper left-hand corner to the lower right-hand corner, making it almost certain that the Red will encounter the Blue on their way. *See Figure 1.* As the Blues traverse the map, they encounter the Red and combat ensues. In each simulation run, if and when the Blues reach their destination, the number of casualties on both sides, and whether the Red Squad reaches its waypoint are noted and entered into our statistics.

In the second experiment, the Blue Squad SA is reduced to 60% and the same number of simulations and observations are repeated for this *non-elite* squad.

The next major step involves splitting the Blues into three squads. The eighteen Blue agents are now distributed into squads of four, six and eight. The addition of two more squads can have profound effects on the way the agents behave as each squad can now decide which of the two tasks (intercepting Red agents or taking their post) to carry out first due to the stochastic nature of the agents' movements. In other words, splitting the Blues into three squads introduces greater complexity than the number may suggest.

First part of our experiment with three Blue squads begins with all squads having the same SA value of 86%. While this new formation of squads consists of the same number of elite agents and same nominal combat capability and situational awareness, the way the tasks are carried out can be entirely different from that in the previous two experiments. For example, one squad could decide to move straight away to take over the Red Post while the other two could choose to stay behind to intercept the Reds. The central question then is whether the *sum* of the *parts* is equal to the *whole*. In other words, are we likely to observe any emergent behaviour?

To complete our investigation into uncertainty in C2 and the relevance of information superiority in C2 and combat performance, we now assign an SA value of 86% to the Blue squad of four, 50% to the squad of six and 33% to the squad of eight. This simulates the effects of each commander interpreting the information differently. Indeed, the randomness or entropy of each squad now truly plays the role of making its own decisions as each squad faces the combined tasks of moving to the next waypoint and intercepting the Reds.

Even though it is unrealistic to expect the Blues to be able to make 100% use of available information from the UAV, two more experiments are conducted in which simulation runs are done on a single Ultra Elite Blue Squad and three Ultra Elite Blue Squads, all with 100% SA. The idea is to use the results as reference points as well as to see if unexpected patterns of behaviour emerge.

Results

The results of the six scenarios are shown in **Charts 1** and **2**. Each set of scenario results is the average of forty simulations runs. While casualties are automatically recorded in a multiple simulation run in the current version of MANA, there is no distinction between the downing of the UAV and the killing a Blue agent. Nor it is possible to automatically record the arrival of a squad at its waypoint (capturing of Red or Blue Post). Therefore, all the simulation runs were watched carefully on the PC monitor screen as the agents moved and fought on the MANA landscape. Though tedious and time-consuming, these direct observations had the benefit of yielding information about the behaviours of the agents, thus helping to understand, for example, why attrition rate is always higher in the 3-squad scenarios than the single-squad ones.

Note that the relatively high probability of the UAV being shot down should not be viewed with too much concern. The UAV is simply a symbolic representation of a source of superior information on the map of MANA. It has very high stealthiness but no firepower. As a result, the Reds tend to shoot at the UAV repeatedly whenever the Blue agents are not shooting at them. While it would have been a simple matter to give the UAV a greater tendency to move away from the Reds, the downing of the

UAV was thought to be another test of the relevance of information superiority to the Blues' mission effectiveness. For example, is there a correlation between the downing of the UAV and the killing of Blue agents?

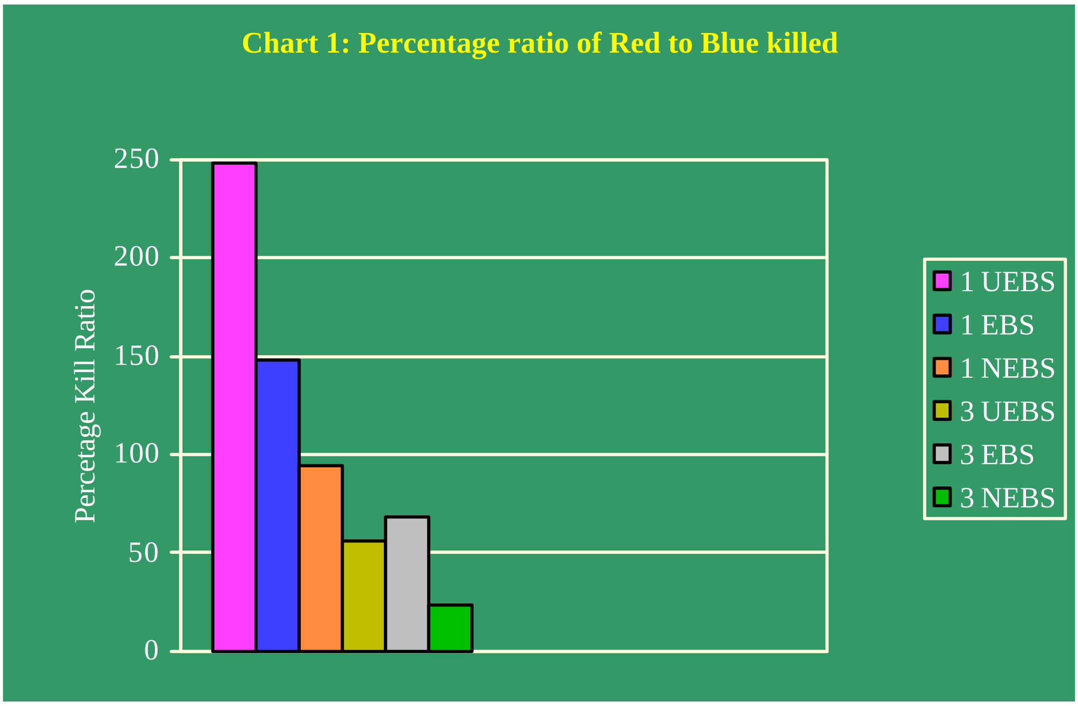

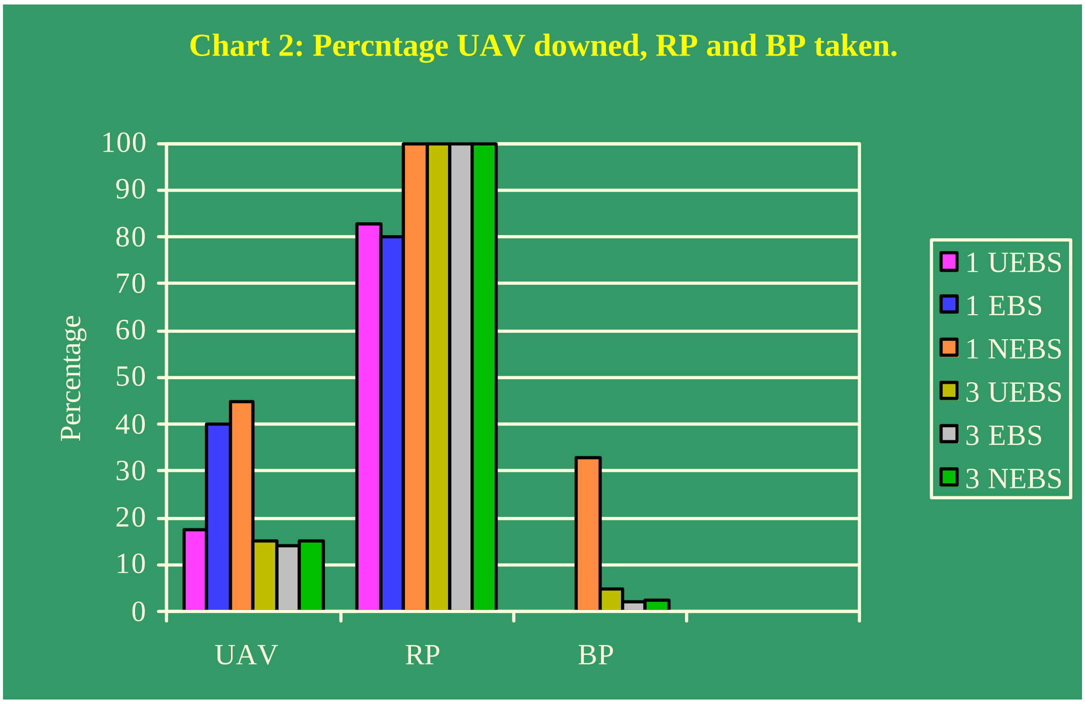

Analysis

The charts show the measures of fitness (MOF) of the Blue agents in each scenario. Recall that the MOF for the Blues include the capturing of the Red Post (Blue's next waypoint) within a fixed time period, intercepting the Red Squad and stopping them from reaching their next waypoint (Blue Post), and avoiding casualties on their side. Although the number of Red agents killed is not one of the tasks, it has a clear connection with the SA value of the Blues. Therefore, we have presented in Chart 1 the percentage ratio of Red to Blue killed as part of our MOF.

It should be stressed once more that the killing of the agents and the capturing of posts are no more than measures of how well the C2 system performs under the influence of uncertainty. It does not necessarily represent a real intercept-combat mission.

Even a cursory look at the charts reveals some rich behavioural pattern of the Blues. First, the percentage kill ratio of the single-squad Blues shows a strongly nonlinear, exponential-like, dependence on the SA value, though one should not immediately conclude that the functional relationship is necessarily exponential based on our rather small sample-parameter space. The reason for the very high percentage kill ratio at the high SA end is easy to understand. Considering that the Blues have superior firepower and weapon range, it is simply a manifestation of a successful first strike. These figures, however, tells us only part of the story. The full picture of the agents' behaviours emerges only from monitoring the computer screen continuously, and thus the reason for the relatively small number of simulation runs.

In the single-squad scenarios, a higher SA value leads to higher kill ratio, lower UAV loss rate and perfect record of preventing Blue Post from being taken. At 100% SA, the Ultra Elite Blue Squad could usually locate the Red agents quickly and kill them before moving on to the Red Post, leaving the Reds little chance to attack the UAV. Note, however, that their ultra high SA value would sometimes make the Blues move very cautiously, and spend a great deal of time digging in and waiting for the Reds to make the next move. This alternative behavioural pattern of the single-squad Blues with high SA values accounts for about 20% of the time when the Blues failed to take the Red Post.

When the single-squad Blues are less well aware of potential hostility, two interesting patterns emerged. First, the Red agents seem to become more cautious and tend to dig in and to try to shoot down the less threatening target of the UAV instead. This is because now both sides have comparable information of each other. Meanwhile, the Blues would take longer to locate the Reds and even longer to neutralise their threat. In some instances, the Blues would completely fail to detect the Reds, and would move to take the Red Post right away. This gives the Reds agents ample opportunity to shoot down the UAV and also the time to take the Blue Post. This behavioural pattern accounts for the sharp increase in UAV being downed and even sharper increase in Blue Post being captured as Blues' SA value drops down to 60%, while at the same time the Red Post capturing rate goes up to 100%.

As mentioned above, there is a great deal more complexity in the behaviour of the agents when the Blues are split into three squads. To begin with, the role of SA has become less transparent in all aspects of MOF, accompanied by the corresponding behavioural change of the agents. The Blues agents now suffer many more casualties than before, while the kill ratio drops dramatically. This is because at most two Blue squads would choose to engage the Reds at any one time, while the other squad would opt for the Red Post. Indeed, the Blue squad with the highest SA would usually engage the enemy first, and would stay to fight, while the other two would choose to move on to their next waypoint. Consequently, the Blues have much smaller concentration of firepower and thus suffers greater attrition loss. On the other hand, the success rate in reaching Red Post is now 100 percent regardless of their situational awareness. With the Blues now getting killed on a regular basis and often not having enough firepower to eliminate the Reds, the Reds can reach their destination with greater frequency.

Two more behavioural patterns were observed in the 3-squad scenarios. First, since the three squads now always split up their tasks, there are at most two Blue squads at any time to engage the Reds. The Red agents tend to become more aggressive, and are now more inclined to shoot at the Blues, as they don't have numerical superiority any longer. Ironically, this often resulted in a greater survival rate for the UAV because it has more time to move out of Red's weapon range while the agents were busy engaged in combat.

Second, the 3-squad Ultra Elite Blues appear to be rather eager to engage the Reds regardless whether the numbers were in their favour or not. For example, the first squad that detects the enemy would engage it without waiting for the arrival and support of another squad. The consequence is that they often suffer heavy loss without achieving the goal of preventing the Reds from capturing the Blue Post.

On a tactical note, the Blues need not kill the Reds to stop them from taking the Blue Post. It is sufficient for some of them to dig in and pose a constant threat to the Reds. The Red agents would not move until they feel no more threat from the Blues. By firing at the Reds too soon, the Blues exposed themselves and resulted in their own elimination. It does seem to confirm that foolish bravery does not pay.

It cannot be over emphasised that models like MANA and ISAAC are stochastic in nature and therefore require a large number of simulations in order to get a reliable indication of the trend. It would also be highly informative to scan over the parameter space in order to study what is called the *fitness landscape*. For example, the MCCDC conducts their study on emergent behaviour in warfare, using ISAAC or similar tools, on a cluster of supercomputers in Maui, Hawaii, for generating very large parameter sample space with the hope of capturing nonlinearity and emergent properties on the fitness landscape. This process is known as *data farming*.

Conclusion

It has been demonstrated that even a simple model like MANA is capable of revealing some surprising patterns of behaviour of interacting agents within a simple command and control structure. It shows that a C2 system is complex and dynamic. Even in our hugely simplified case, there are surprises and even nonlinear features emerging from the interactions among the agents. It clearly shows that a small increase in C2 complexity (from a single-squad operation to a 3-squad one) can lead to significant increase in randomness (uncertainty) and behavioural change. Such features could not have been captured using physics-based models. The so-called New Sciences of complexity theory offer new methodologies and tools for investigating complex problems in military operations and systems, such as command and control, and may provide new insights into their workings.

This work may also contribute to understanding the role of information superiority in C2 effectiveness. The availability of high quality information is no good in the hands of poor commanders who either do not understand it or know how to use it properly. It is also clear that if superior information is not used intelligently, with appropriate adjustment in tactics and even doctrine, it may actually be detrimental to mission effectiveness (see the results of 3-squad Ultra Elite Blues and the discussions).

This work is a preliminary study of applying the relatively new tool of agent-based models to the century old C2 problems. The results presented here are based on "toy models" [4], which are not very sophisticated in several features. As stated repeatedly before, a much larger sample-parameter space would be required before true emergent behaviours and nonlinearities are captured. Last but not least, the availability of only two levels of command in the current version of MANA is unsatisfactory for a realistic study of the functioning of C2. The next stage of our work will aim to overcome some of the difficulties encountered in using these models.

Acknowledgment

Dr. Michael Lauren of DOTSE, New Zealand, kindly provided us with a copy of the MANA software, which made this work possible in a very short space of time. There have also been many useful discussions between Dr. Lauren and the author. All these are gratefully acknowledged. The author would like to thank DSTO for financial support to attend the *3rd Project Albert International Workshop* in Auckland, New Zealand, 19-22 February 2001. This Workshop has proved to be immensely useful for getting this project off the ground, and for establishing new collaborations.

References

1. *Manoeuvre Warfare Science 2001*, ed. Gary Horne and Mary Leonardi, Marine Corps Combat Development Command, Copyright © 2001.
2. *Complexity, Global Politics, and National Security*, ed. D.S. Alberts and T.J. Czerwinski, National Defense University, Washington. http://www.dodccrp.org.
3. John F. Schmitt, Command and (Out of) Control: The Military Implication of Complexity Theory, *ibid* Chapter 9.
4. Andrew Ilachinski, *Towards a Science of Experimental Complexity: An Artificial-Life Approach to Modelling Warfare*, http://www.can.org.isaac/.
5. *Command in War*, Martin van Creveld, Harvard University Press (1985).
6. Andrew Ilachinski, *Irreducible Semi-Autonomous Adaptive Combat (ISAAC): An Artificial-Life Approach to Land Combat*, Military Operations Research, **5**, No.3 29.
7. Michael K. Lauren, *Exploring the Value of Sensors to a Recce Unit Using Agent-based Models*, J. Battlefield Technology, **4**, No.1, March 2001.
8. Michael K. Lauren and Roger T. Stephen, *Modelling Patrol Survivability in A Generic Peacekeeping Setting Using ISAAC*, DOTSE Report **177**, NR 1358, ISSN 1174-3387.